

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

Appendix A: Claims as Currently Presented

Claims 1 - 24 have been canceled.

1 Claim 25: A method of searching and updating indexes to a data structure in a multi-processing
2 environment, comprising steps of:
3 maintaining two indexes to a single data structure, a first index for searching and a second
4 index for updating;
5 responsive to each update of the second index, switching the indexes so that the first index
6 becomes the second index and the updated second index becomes the first index;
7 allowing searches that are in progress using the first index, before the switching, to
8 continue until completion after the switching, using the newly-switched second index;
9 after the switching, initiating new searches using the newly-switched first index;
10 when all searches in the newly-switched second index have completed, updating the
11 newly-switched second index in an identical manner as the update to which the switching step was
12 responsive; and
13 preventing another operation of the switching step until completion of the step of updating
14 the second index in the identical manner.

Claim 26 has been canceled.

1 Claim 27: A program product storage medium containing computer instructions that when
2 executed in a computer perform a method of searching and updating indexes to a data structure in

3 a multi-processing environment, the method comprising steps of:

4 maintaining two indexes to a single data structure, a first index for searching and a second
5 index for updating;

6 responsive to each update of the second index, switching the indexes so that the first index
7 becomes the second index and the updated second index becomes the first index;

8 allowing searches that are in progress using the first index, before the switching, to
9 continue until completion after the switching, using the newly-switched second index;

10 after the switching, initiating new searches using the newly-switched first index;

11 when all searches in the newly-switched second index have completed, updating the
12 newly-switched second index in an identical manner as the update to which the switching step was
13 responsive; and

14 preventing another operation of the switching step until completion of the step of updating
15 the second index in the identical manner.

Claim 28 has been canceled.

1 Claim 29: A computer program product for serializing data structure retrievals and updates in a
2 multi-processing computer system, the computer program product embodied on one or more
3 computer-readable media and comprising:

4 computer-readable program code means for creating two identical data structures, both
5 representing an initial state for accessing a single copy of stored data;

6 computer-readable program code means for performing searches against a first of the two

7 data structures, the computer-readable program code means for performing searches further
8 comprising a first program instruction for incrementing a search use count for the first data
9 structure atomically during each search to ensure no interference from other processes during that
10 search and a second instruction for decrementing the search use count for the first data structure
11 atomically after performing each search;

12 computer-readable program code means for performing a first update against a second of
13 the two data structures, yielding a revised data structure;

14 computer-readable program code means for switching the first data structure and the
15 revised data structure, responsive to completion of the computer-readable program code means
16 for performing the first update, such that the first data structure becomes the second data
17 structure and the revised data structure becomes the first data structure, the computer-readable
18 program code means for switching the data structures further comprising a third instruction for
19 re-ordering data structure pointers atomically to prevent interference from other processes during
20 operation of the computer-readable program code means for switching; and

21 computer-readable program code means for applying, after operation of the computer-
22 readable program code means for switching, the first update against the second data structure,
23 yielding a second data structure that is structurally identical to the first data structure;

24 the computer-readable program code means for performing searches further comprising
25 computer-readable program code means for activating the computer-readable program code
26 means for applying the first update against the second data structure when the search use count
27 for the second data structure has a value indicating that no searches are being performed against
28 the second data structure.

1 Claim 30: The computer program product according to Claim 29, further comprising:

2 computer-readable program code means for obtaining an exclusive lock on the second
3 data structure prior to operation of the computer-readable program code means for performing
4 the first update; and

5 computer-readable program code means for releasing the exclusive lock after operation of
6 the computer-readable program code means for applying the first update.

1 Claim 31: The computer program product according to Claim 29, wherein the computer-readable
2 program code means for performing the first update further comprises computer-readable
3 program code means for queuing a transaction that specifies one or more data structure traversals
4 and one or more data structure modifications that were performed to yield the revised data
5 structure, and wherein the computer-readable program code means for applying the first update
6 further comprises computer-readable program code means for performing the one or more data
7 structure traversals and the one or more modifications specified in the queued transaction against
8 the second data structure that results from operation of the computer-readable program code
9 means for switching.

1 Claim 32: The computer program product according to Claim 29, further comprising computer-
2 readable program code means for performing a subsequent update against the second data
3 structure that results from operation of the computer-readable program code means for applying
4 the first update; and wherein operation of the computer-readable program code means for

5 performing the subsequent update causes another operation of the computer-readable program
6 code means for switching and the computer-readable program code means for applying.

1 Claim 33: A computer system for serializing data structure retrievals and updates in a multi-
2 processing computer system, the computer system comprising:

3 means for creating two identical data structures, both representing an initial state for
4 accessing a single copy of stored data;

5 means for performing searches against a first of the two data structures, the means for
6 performing searches further comprising means for incrementing a search use count for the first
7 data structure atomically during each search to ensure no interference from other processes
8 during that search and means for atomically decrementing the search use count for the first data
9 structure after performing each search;

10 means for performing a first update against a second of the two data structures, yielding a
11 revised data structure;

12 means for switching the first data structure and the revised data structure, responsive to
13 completion of the means for performing the first update, such that the first data structure becomes
14 the second data structure and the revised data structure becomes the first data structure, the
15 means for switching the data structures further comprising means for re-ordering data structure
16 pointers atomically to prevent interference from other processes during operation of the means
17 for switching; and

18 means for applying, after switching the data structures, the first update against the second
19 data structure, yielding a second data structure that is structurally identical to the first data

20 structure;
21 the means for performing searches further comprising means for activating the means for
22 applying the first update against the second data structure when the search use count for the
23 second data structure has a value indicating that no searches are being performed against the
24 second data structure.

1 Claim 34: The system according to Claim 33, further comprising:
2 means for obtaining an exclusive lock on the second data structure prior to operation of
3 the means for performing the first update; and
4 means for releasing the exclusive lock after operation of the means for applying the first
5 update.

1 Claim 35: The system according to Claim 33, wherein the means for performing the first update
2 further comprises means for queuing a transaction that specifies one or more data structure
3 traversals and one or more data structure modifications that were performed to yield the revised
4 data structure, and wherein the means for applying the first update further comprises means for
5 performing the one or more data structure traversals and the one or more data structure
6 modifications specified in the queued transaction against the second data structure that results
7 from operation of the means for switching.

1 Claim 36: The system according to Claim 33, further comprising means for performing a
2 subsequent update against the second data structure that results from operation of the means for

3 applying the first update; and wherein operation of the means for performing the subsequent
4 update causes another operation of the means for switching and the means for applying.

1 Claim 37: A method for serializing data structure retrievals and updates in a multi-processing
2 computer system, comprising steps of:

3 creating two identical data structures, both representing an initial state for accessing a
4 single copy of stored data;

5 performing searches against a first of the two data structures, the performing searches step
6 further comprising the step of incrementing a search use count for the first data structure
7 atomically during each search to ensure no interference from other processes during the search
8 and the step of decrementing the search use count for the first data structure atomically after
9 performing each search;

10 performing a first update against a second of the two data structures, yielding a revised
11 data structure;

12 switching the first data structure and the revised data structure, responsive to completion
13 of the step of performing the first update, such that the first data structure becomes the second
14 data structure and the revised data structure becomes the first data structure, the step of
15 switching the data structures further comprising the step of re-ordering data structure pointers
16 atomically to prevent interference from other processes during operation of the switching step;
17 and

18 applying, after the switching step, the first update against the second data structure,
19 yielding a second data structure that is structurally identical to the first data structure;

20 the step of performing searches further comprising the step of activating the step of
21 applying the first update against the second data structure when the search use count for the
22 second data structure has a value indicating that no searches are being performed against the
23 second data structure.

1 Claim 38: The method according to Claim 37, further comprising steps of:
2 obtaining an exclusive lock on the second data structure prior to performing the first
3 update; and
4 releasing the exclusive lock after applying the first update.

1 Claim 39: The method according to Claim 37, wherein the step of performing the first update
2 further comprises the step of queuing a transaction that specifies one or more data structure
3 traversals and one or more data structure modifications that were performed to yield the revised
4 data structure, and wherein the step of applying the first update further comprises the step of
5 performing the one or more data structure traversals and the one or more data structure
6 modifications specified in the queued transaction against the second data structure that results
7 from operation of the switching step.

1 Claim 40: The method according to Claim 37, further comprising the step of performing a
2 subsequent update against the second data structure that results from applying the first update;
3 and wherein the step of performing the subsequent update causes repeating the switching step and
4 the applying step.

1 Claim 41: A method for serializing data retrievals and updates in a computing environment,
2 comprising steps of:

3 creating two identical indexes, both representing an initial state for accessing stored data
4 and each indexing a single copy of the stored data;

5 performing searches against a first of the two indexes;

6 performing a first update against a second of the two indexes, yielding a revised index;

7 serializing information describing a traversal path taken through the second index for
8 making the first update and one or more modifications made to the second index in the first
9 update;

10 switching the first index and the revised index, responsive to performing the first update,
11 such that the first index becomes the second index and the revised index becomes the first index;

12 applying, after the switching step, the first update to the second index, using the serialized
13 information describing the traversal path and the one or more modifications to traverse and
14 modify the newly-switched second index, thereby yielding a second index that is synchronized
15 with, and structurally identical to, the first index; and

16 performing subsequent searches against the first index.

1 Claim 42: The method according to Claim 41, further comprising the step of performing a
2 subsequent update against the second index that results from applying the first update; and
3 wherein the step of performing the subsequent update causes repeating the serializing, switching,
4 and applying steps.

1 Claim 43: A method of serializing access to data in a computing system, comprising steps of:

2 maintaining two trees as indexes to a single copy of data, a first of which is used for
3 searches and a second of which is used for update operations, each tree having a use count
4 associated therewith;

5 carrying out searches using the search tree, further comprising the steps of:

6 determining, for each new search request, which of the trees is currently the search
7 tree;

8 incrementing the use count for the search tree;

9 performing the new search request using the search tree; and

10 decrementing the use count for the search tree, responsive to completion of the
11 performing step; and

12 carrying out each update using the update tree, further comprising the steps of:

13 determining which of the trees is currently the update tree;

14 performing an update to the update tree;

15 serializing a record describing the update to the update tree;

16 switching the update tree to become the search tree and the search tree to become
17 the update tree, responsive to completion of the steps of performing the update and serializing the
18 record; and

19 applying the serialized record to the newly-switched update tree, provided that the
20 use count for the newly-switched update tree has reached a value that indicates that no search
21 requests are currently being performed against this newly-switched update tree, delaying the step

22 of applying the serialized record if necessary until the use count for the newly-switched update
23 tree has reached this value, and wherein the step of applying the serialized record ensures that
24 both the search tree and the update tree reflect each update.

1 Claim 44: The method according to Claim 41, wherein the indexes are implemented as trees.

1 Claim 45: The method according to Claim 41, wherein the indexes are implemented as hash
2 tables.